



U.S. DEPARTMENT OF
ENERGY

Advanced Scientific Computing Research Research Priorities Overview

Karen Pao

Advanced Scientific Computing Research (ASCR)

Karen.Pao@science.doe.gov

5 January 2011



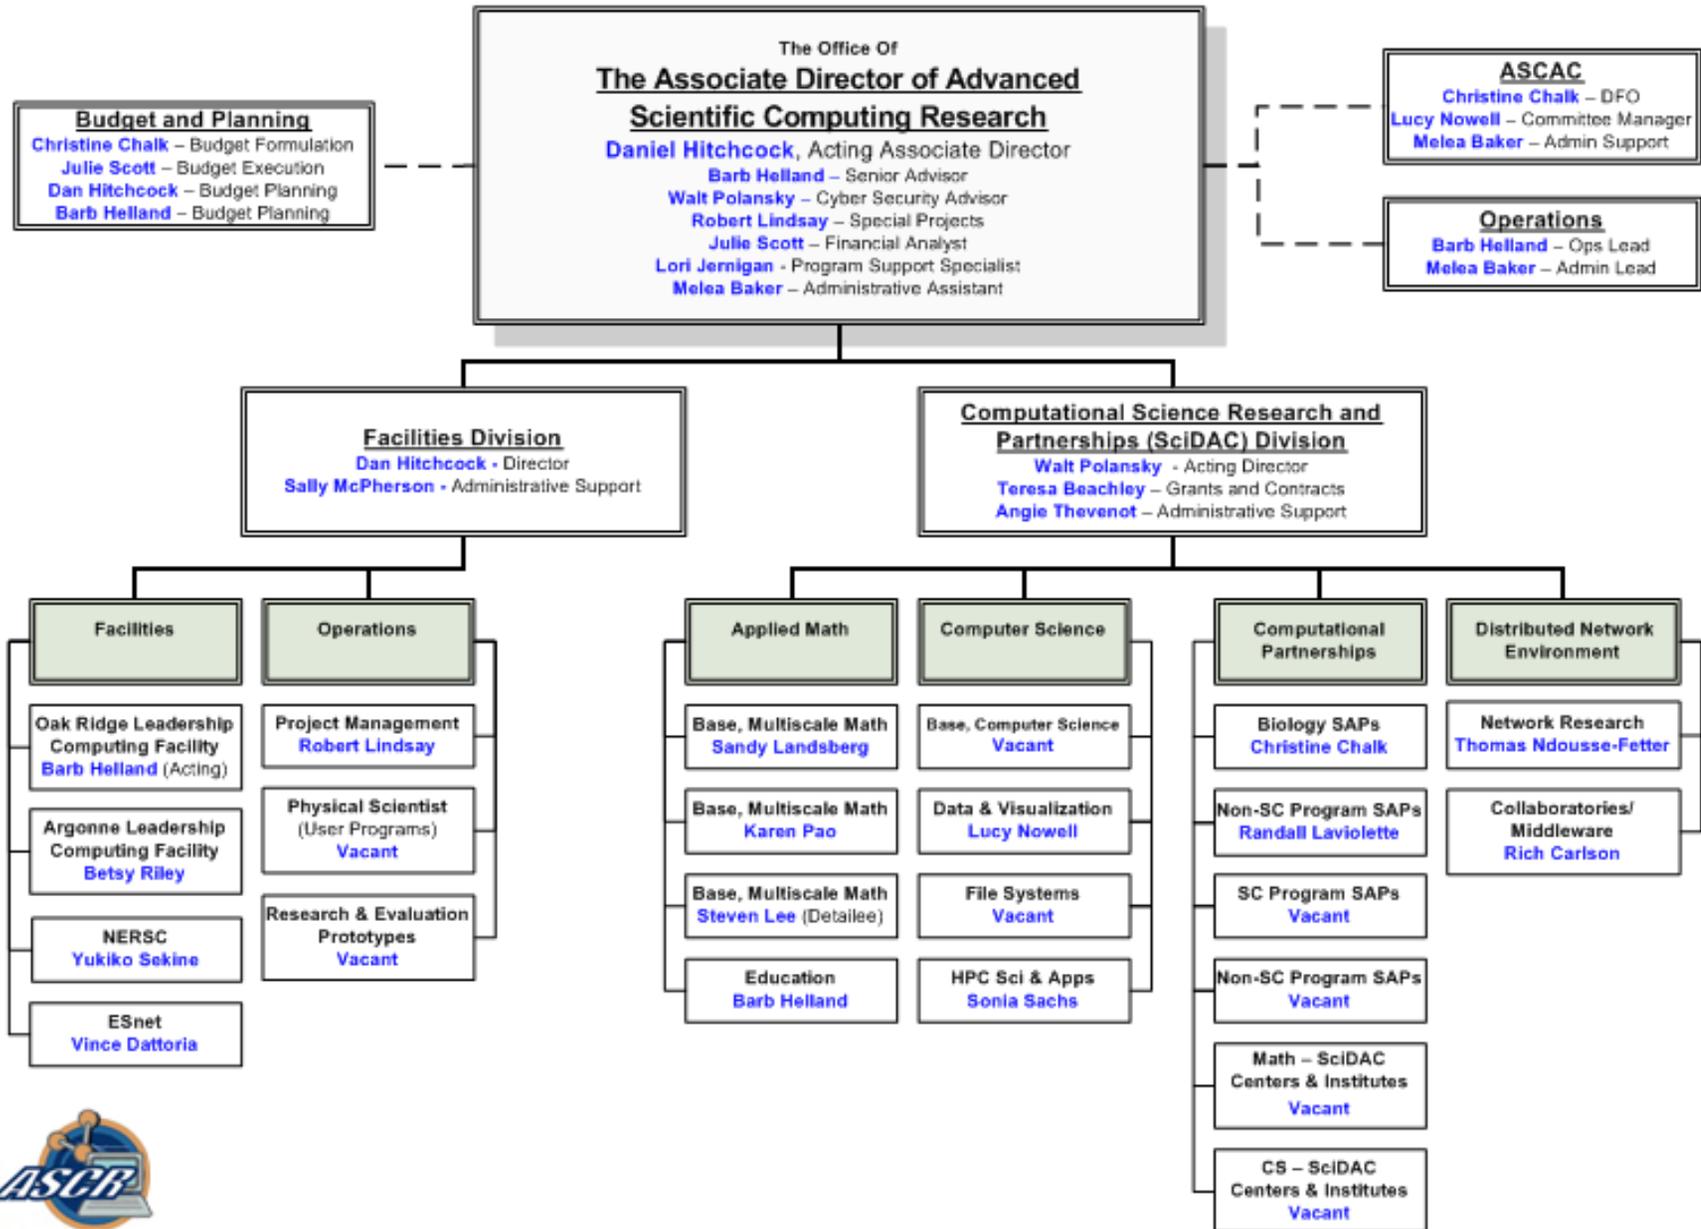
ASCR Mission

The mission of the Advanced Scientific Computing Research (ASCR) program is to *discover, develop, and deploy the computational and networking capabilities* that enable researchers to analyze, model, simulate, and predict complex phenomena important to the Department of Energy.

A particular challenge of this program is *fulfilling the science potential of emerging multi-core computing systems and other novel “extreme-scale” computing architectures*, which will require significant modifications to today’s tools and techniques to deliver on the promise of exascale science.



THE OFFICE OF
ADVANCED SCIENTIFIC COMPUTING RESEARCH
 Functional Organization Chart



Exascale Computing

Meeting DOE Secretarial Priorities

Supporting Exascale program will meet DOE's Secretarial Priorities by

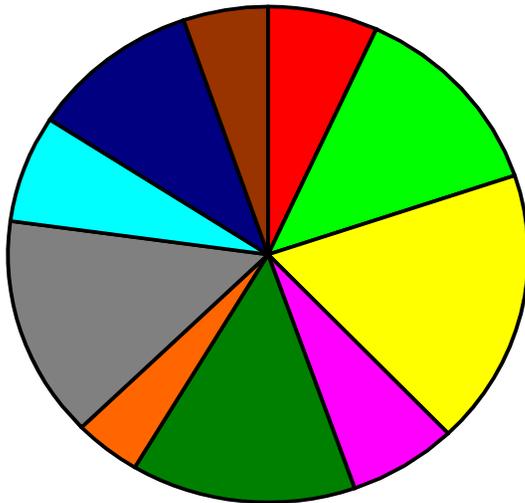
- Improving predictive understanding in complex systems (e.g. **Climate, Energy, Combustion, Materials and Nuclear Weapons Performance**), including coupling of uncertainty quantification with exascale systems simulations
- Achieving the power efficiency, reliability and programmability goals essential for exascale -- enormous positive impact on business information technology as well as mid-range scientific computing
- Providing a ten year goal to drive modeling and simulation developments forward and to attract young scientific talent to DOE

By investing now, the US will achieve an exascale simulation capability within ten years and doing so will enable the US to lead in key areas of science and engineering

- Exascale will be here by 2020...the next few years will be interesting!
- We have no idea what it will look like – but we have the two “Swim Lanes” – two current design points for achieving the performance required by science and engineering applications – that may still change
- “Co-Design”, or trying to get the software ready as the hardware is being developed, has not been practiced before -- it has been “here’s the next machine; now you figure out how to use it” for the last 20 years.
- We need solutions to handle decreased reliability and a new model for resiliency
- We know that system software as currently implemented is not suitable for exascale system
- With more expensive machines, there will be demand for *demonstrably* higher-quality science (not just prettier pictures!)
- With the explosion of data, it is quite possible that different kinds of application may dominate NERSC usage

***All these make it rather difficult to predict
future requirements!***

Research Supported in 2010



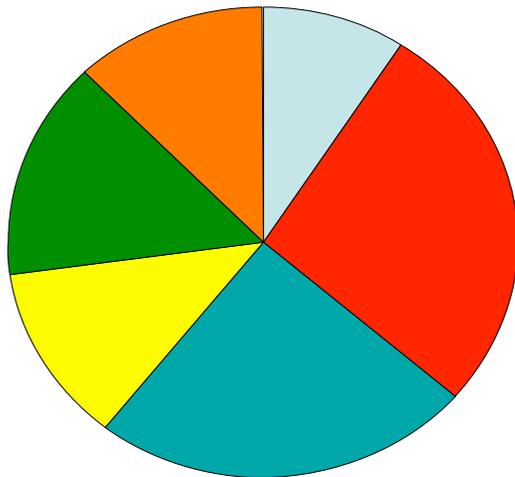
- | | |
|-----------------------------------|--------------------------------|
| ■ Advanced Linear Algebra | ■ Discretization & Meshing |
| ■ Multiscale, Multiphysics | ■ Uncertainty & Error Analysis |
| ■ Optimization | ■ Other research |
| ■ Fellowships & Workshops | ■ Cross-Cut Centers |
| ■ Math Analysis of Petascale Data | ■ Joint Math-CS Institutes |

- **Uncertainty Quantification**
- **Numerical methods for PDE**
- **Advanced linear algebra**
- **Optimization algorithms**
- **Joint Applied Mathematics-Computer Science Institutes**
- **Mathematics for the analysis of extremely large datasets**
- **Mathematics of cyber security basic research**
- **Computational Science Graduate Fellowship**



Computer Science

Research supported in 2010



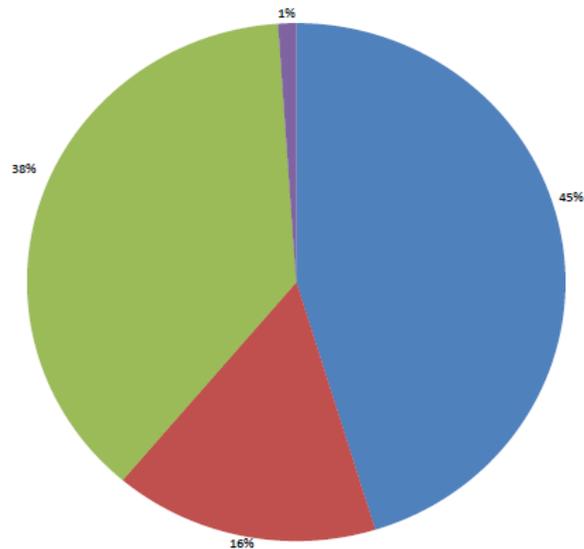
- Data Management and Visualization
- Programming Models
- Operating and File Systems
- Math CS Institutes
- Performance and Productivity Tools
- Workshops/Other

- Extreme Scale Operating and file systems.
- Performance and productivity tools
- Programming models
- Data management and visualization
- Joint Applied Mathematics-Computer Science Institutes
- Pre-competitive research in advanced computer architectures



Computational Partnerships

Research supported in 2010

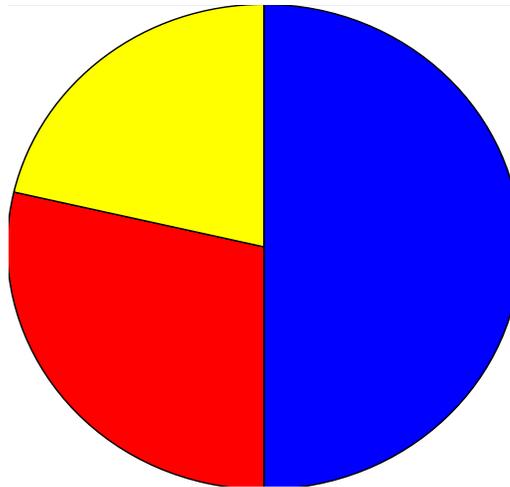


- Centers for Enabling Technologies
- SciDAC Institutes
- Science Applications and Partnerships
- SciDAC Outreach Center

- **SciDAC Institutes**
- **Scientific Computation Application Partnerships**
- Department of Energy Computation Application Partnerships (e.g. Nuclear Energy and/or Office of Electricity Delivery and Energy Reliability)
- Co-Design Centers will continue the close coupling of applications, computer science, and computer hardware architecture that is required for success at exascale.

Next Generation Networking for Science

Research supported in 2010



- Collaboratories and Middleware
- Advanced Network Services
- Workshops/Other

- Collaboration tools for grand-challenge scientific research.
- Distributed systems software including scalable and secure tools and services to facilitate large-scale national and international scientific collaborations
- Advanced network technologies including dynamic optical network services, scalable cyber security technologies, and multi-domain, multi-architecture performance protocols to seamlessly interconnect and provide access to distributed computing resources and science facilities.



U.S. DEPARTMENT OF
ENERGY

?



- **The amount of data (observational, experimental, as well as 0's and 1's from simulations) will be unprecedented**
 - **Example: Fusion:** “Some [fusion codes] will **model 1 billion cells and 1 trillion particles**. Based on mean-time-between-failure concerns when running on a million cores, these codes will need to **output 2 gigabytes/second per core or 2 petabytes/second of checkpoint data every 10 minutes**. This amounts to an unprecedented **input/output rate of 3.5 terabytes/second**. The data questions to consider at the extreme scale fall into two main categories: data generated and collected during the production phase, and data that need to be accessed during the analysis phase.”
 - **Example: Climate:** “Climate model data sets are growing faster than the data set size for any other field of science [...] Based on current growth rates, these data sets would be **hundreds of exabytes** by 2020. To provide the international climate community with convenient access to these data and to maximize scientific productivity, these data would need to be replicated and cached at multiple locations around the globe.”
- **Need to refine and develop methods and technologies to**
 - **Move the data**
 - **Store the data**
 - **Understand the data**



- **Applied Mathematics Research**

- Improve methods for data and dimension reduction are needed to extract pertinent subsets, features of interest, or low-dimensional patterns, from large raw data sets
- Understand uncertainty, especially in messy and incomplete data sets
- Identify, in real time, anomalies in streaming and evolving data is needed in order to detect and respond to phenomena that are either short-lived or urgent

- **Computer Science Research**

- Extreme-scale data storage and access systems for scientific computing, that minimize the need for scientists to have detailed knowledge of system hardware and operating systems;
- Scalable data triage, summarization, and analysis methods and tools for in-situ data reduction and/or analysis of massive multivariate data sets;
- Semantic integration of heterogeneous scientific data sets;
- Data mining, automated machine reasoning, and knowledge representation methods and tools that support automated analysis and integration of large scientific data sets, especially those that include tensor flow fields; and
- Visual analysis of extreme-scale scientific data, including multi-user visual analysis methods and tools and interactive visual steering of computational processes.

- **Next-generation Networking Research (workshop in the works)**

- Deploy high-speed networks for effective and easy data transport
- Develop real-time network monitoring tools to maximize throughput
- Manage collections of extreme scale data across a distributed network

Requirements for NERSC

- **Scientific data management (Shoshani)**
- **Future data analysis needs (Samatova)**
- **I/O software (Choudhary)**
- **Visualization (Ma)**
- **Visualization (Bethel)**

- The overarching goal of SciDAC has always been to “*bring together the nation's top researchers to tackle challenging scientific problems by advancing computational science and developing the tools necessary to enable scientific use of the high performance computers of today and those envisioned for the next decade.*”
- Some of the biggest users of ASCR NERSC allocations today are “CFD-like applications” or “PDE-based applications”
- Major computational challenge of these applications today: scalability
 - NERSC as a “springboard” for ALCF

Exascale Computing brings new challenges to applications

- **Traditionally, PDE-based applications have expected 10x increase in resolution with each 1000x increase in compute capability, *but not this time:***
 - We won't have 1000x the memory available
 - The processors won't be 10x faster
 - Proportionally, we won't be able to move as much data on or off each processor
 - Introduction of massive parallelism at the node level is a significant new challenge (MPI is only part of the solution)
- **However, exascale computing is an opportunity for...**
 - ***More Fidelity:*** Incorporate more physics instead of increased resolution
 - ***Greater Understanding:*** Develop uncertainty quantification (**UQ**) to establish confidence levels in computed results and deliver predictive science
- **Memory management will be key in PDE applications**
 - Recast applied math algorithms/PDEs discretization to reflect shift from FLOP-centric to memory-constrained hardware
 - New algorithms with more compute, less communication
- **Applications must care more about fault tolerance and resilience**
 - Checkpoint-Restart may not work with current storage & I/O systems



Uncertainty Quantification (UQ)

- **The need for UQ in nuclear weapons application simulations became more obvious as experiments become harder and more expensive to conduct**
 - Simulation and modeling become the main tools for “stockpile stewardship”
 - Old experiments are reanalyzed and error bars on data are (re)established
 - Most experimental data are “integral” in nature and in general do not validate any individual physics model in the multi-physics context
 - Simulation and modeling play an important role in designing necessary (above-ground) experiments
- **The need for UQ for the climate modeling community has also become clear**
 - Like the NW community, no full-up integral experiment can be conducted
 - Like the NW community, many uncertainties exist in individual “physics” models and the data can only “validate” integrated systems
 - Like the NW community, policy-makers need to understand the uncertainties in the simulations as they decide on a course of action (e.g., investments in low-lying coastal areas, energy policies, etc)
- **The need for UQ for scientific discovery in general isn’t there yet, but as the investments for resources become more expensive, it’s a matter of time before someone asks for it!**

- **UQ is on the critical path of ASCR's March to Exascale**
 - In 2010, ASCR made a modest investment in mathematical methods of uncertainty quantification
 - The ASC program has been making substantial investments in V&V and UQ through its Predictive Science Academic Alliance (PSAAP) program
- **Even though today there's little UQ activities on ASCR facilities, it is anticipated that these activities will increase over time**
 - Climate comes to mind...
 - Nuclear Energy (NE), although not in Office of Science, has also put much emphasis in UQ in its M&S efforts
- **In Today's UQ Case Study we are going to look at an example from a programmatic application at LLNL as a reference point on possible requirements for science applications**



NERSC Requirements for Applications and UQ

- **Simulation and Analysis of Reacting Flows (Bell)**
- **Shock, turbulence, and material interface (Lele)**
- **UQ (Tong)**
- **Not presenting today – APDEC (Colella)**

- **Mathematical software represents the realization of abstract algorithms into something that scientists and engineers – who are often not conversant in latest advances in mathematics – can deploy**
 - Traditionally, these software are used in PDE-based applications
 - Computational chemistry/materials science applications *at larger scales* are gaining importance
- **Math software lies in the heart of every science and engineering application; it is paramount that these software are of the highest quality, embody the most advanced numerical algorithms possible, and are useful and usable by application scientists!**
 - This is where mathematics and computer science intersect: it is not enough to simply provide good mathematics algorithms, but the software needs to run on today's and tomorrow's architectures – and computer science issues such as programming models, data movement, I/O, storage, systems software, etc. all come into play
- **We have numerous NERSC users in this area; most have small (< 100,000 proc hrs) allocations but tend to use large number of processors**
 - Scalability in linear algebra software is today's major challenge
 - and will be tomorrow's as well.

Exascale computing will only add to the complexity of these software

- **Will the codes supported by DOE survive the paradigm shift?**
 - Will codes need to be rewritten?
 - How to make this transition less painful?
 - How to make these codes “architecture aware”?
- **Memory management will again be the key**
 - Need to reduce global communication in linear algebra (!)
 - What to do with those large, dense matrices?
 - How to achieve scalability?
- **Math and CS issues are intertwined:**
 - **Mathematics issues:**
 - Reformulation of existing algorithms to swap flops for data locality
 - Exploration of novel algorithms, such as mixed precision, etc
 - **Computer Science Issues:**
 - Power consumption as a function of algorithms?
 - Fault tolerance/resilience
 - Programming models – how to phase them in

- **Osni Marques -- Installation, Testing and Evaluation of ACTS Tools**
- **Esmond Ng -- High Performance Sparse Matrix Algorithms**



U.S. DEPARTMENT OF

ENERGY

Computer Science & Performance Evaluation

- **Computer science, including performance evaluation and engineering, will play a crucial role in exascale computing:**
 - Exascale computer systems will be high-concurrency and a change from current architecture
 - data movement, rather than processors and computational operations, will be the limiting factor for achieving exascale
 - power constraints will dictate architecture of the exascale system
 - memory per core is expected to decline sharply
 - the performance of storage systems continues to lag far behind; multi-level storage architectures that span multiple types of hardware are anticipated
- **Critical challenges include**
 - Communication
 - Energy efficiency and power management
 - Fault tolerance and resiliency
 - Data locality
 - Memory hierarchy
 - Resource scheduling
- **Computer scientists will conduct research and design tools so that the rest of us would not have to manage these challenges!**

Exascale technology needs go beyond performance evaluation

- **System power** is a first class constraint on exascale system performance and effectiveness.
- **Memory** is an important component of meeting exascale power and applications goals.
- **Programming model.** Early investment in several efforts to decide in 2013 on exascale programming model, allowing exemplar applications effective access to 2015 system for both mission and science.
- **Investment in exascale processor design** to achieve an exascale-like system in 2015.
- **Operating System strategy for exascale** is critical for node performance at scale and for efficient support of new programming models and run time systems.
- **Reliability and resiliency are critical at this** scale and require applications neutral movement of the file system (for check pointing, in particular) closer to the running apps.
- **HPC co-design strategy and implementation** requires a set of a hierarchical performance models and simulators as well as commitment from apps, software and architecture communities.

- **Current ASCR investments include**
 - **Joint Math/CS Institute**
 - **SciDAC projects PERI & CsCADs**
 - **X-Stack**
 - Not included here are the numerous data management/analytics/analysis projects
- **Current there are numerous NERSC users in these areas; these users tend to have small allocations but use many processors (16K and beyond)**



Computer Science Requirements

- **Erich Strohmaier – a compendium of computer science projects**



U.S. DEPARTMENT OF
ENERGY

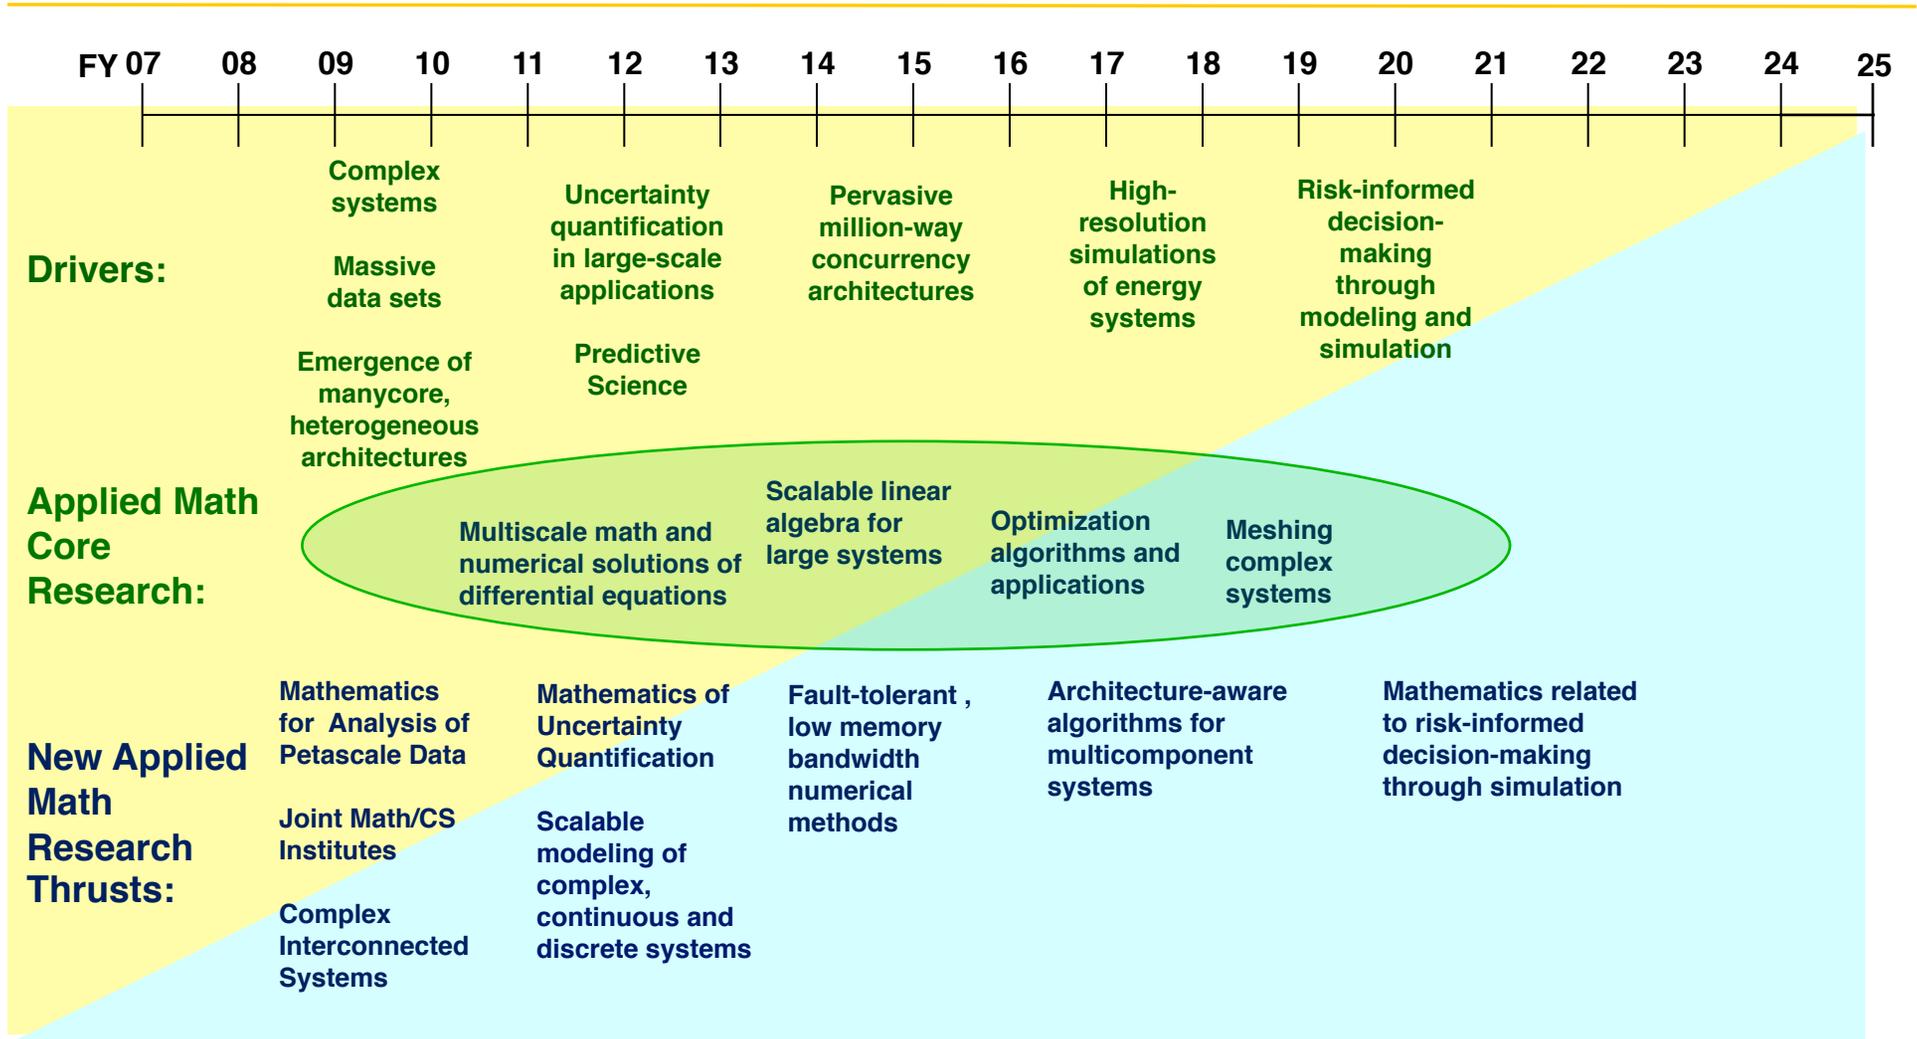
Backup

Progress toward Exascale since May, 2010

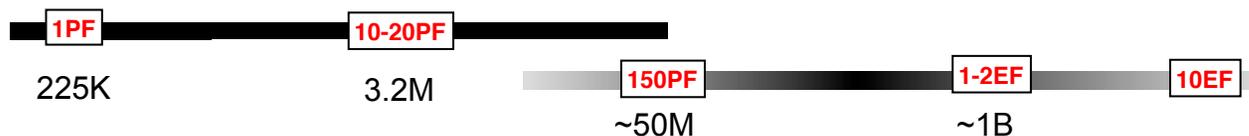
- Proposals processed in Exascale related topic areas:
 - Applied Math: Uncertainty Quantification (90 proposals requesting ~\$45M/year; 6 funded at \$3M/yr)
 - Computer Science: Advanced Architectures(28 proposals requesting ~\$28M/year, 6 funded at \$5M/yr)
 - Computer Science: X-Stack (55 proposals requesting ~\$40M/year, 11 funded at \$8.5M/yr)
 - Computational Partnerships: Co-Design (21 Proposals requesting ~ \$160M/year)
- Exascale Coordination meeting with DOD and DARPA – June 15 with decisions to --
 - Have yearly coordination/planning meetings in accordance with current MOU
 - Follow-on meeting of Program Managers to identify gaps and overlaps in computer science and applied math areas currently funded
 - Address critical technologies issues, initially in memory
- Partnership with NNSA
 - MOU between SC and NNSA in progress
 - Ongoing meetings between ASCR and ASC



Applied Math Research

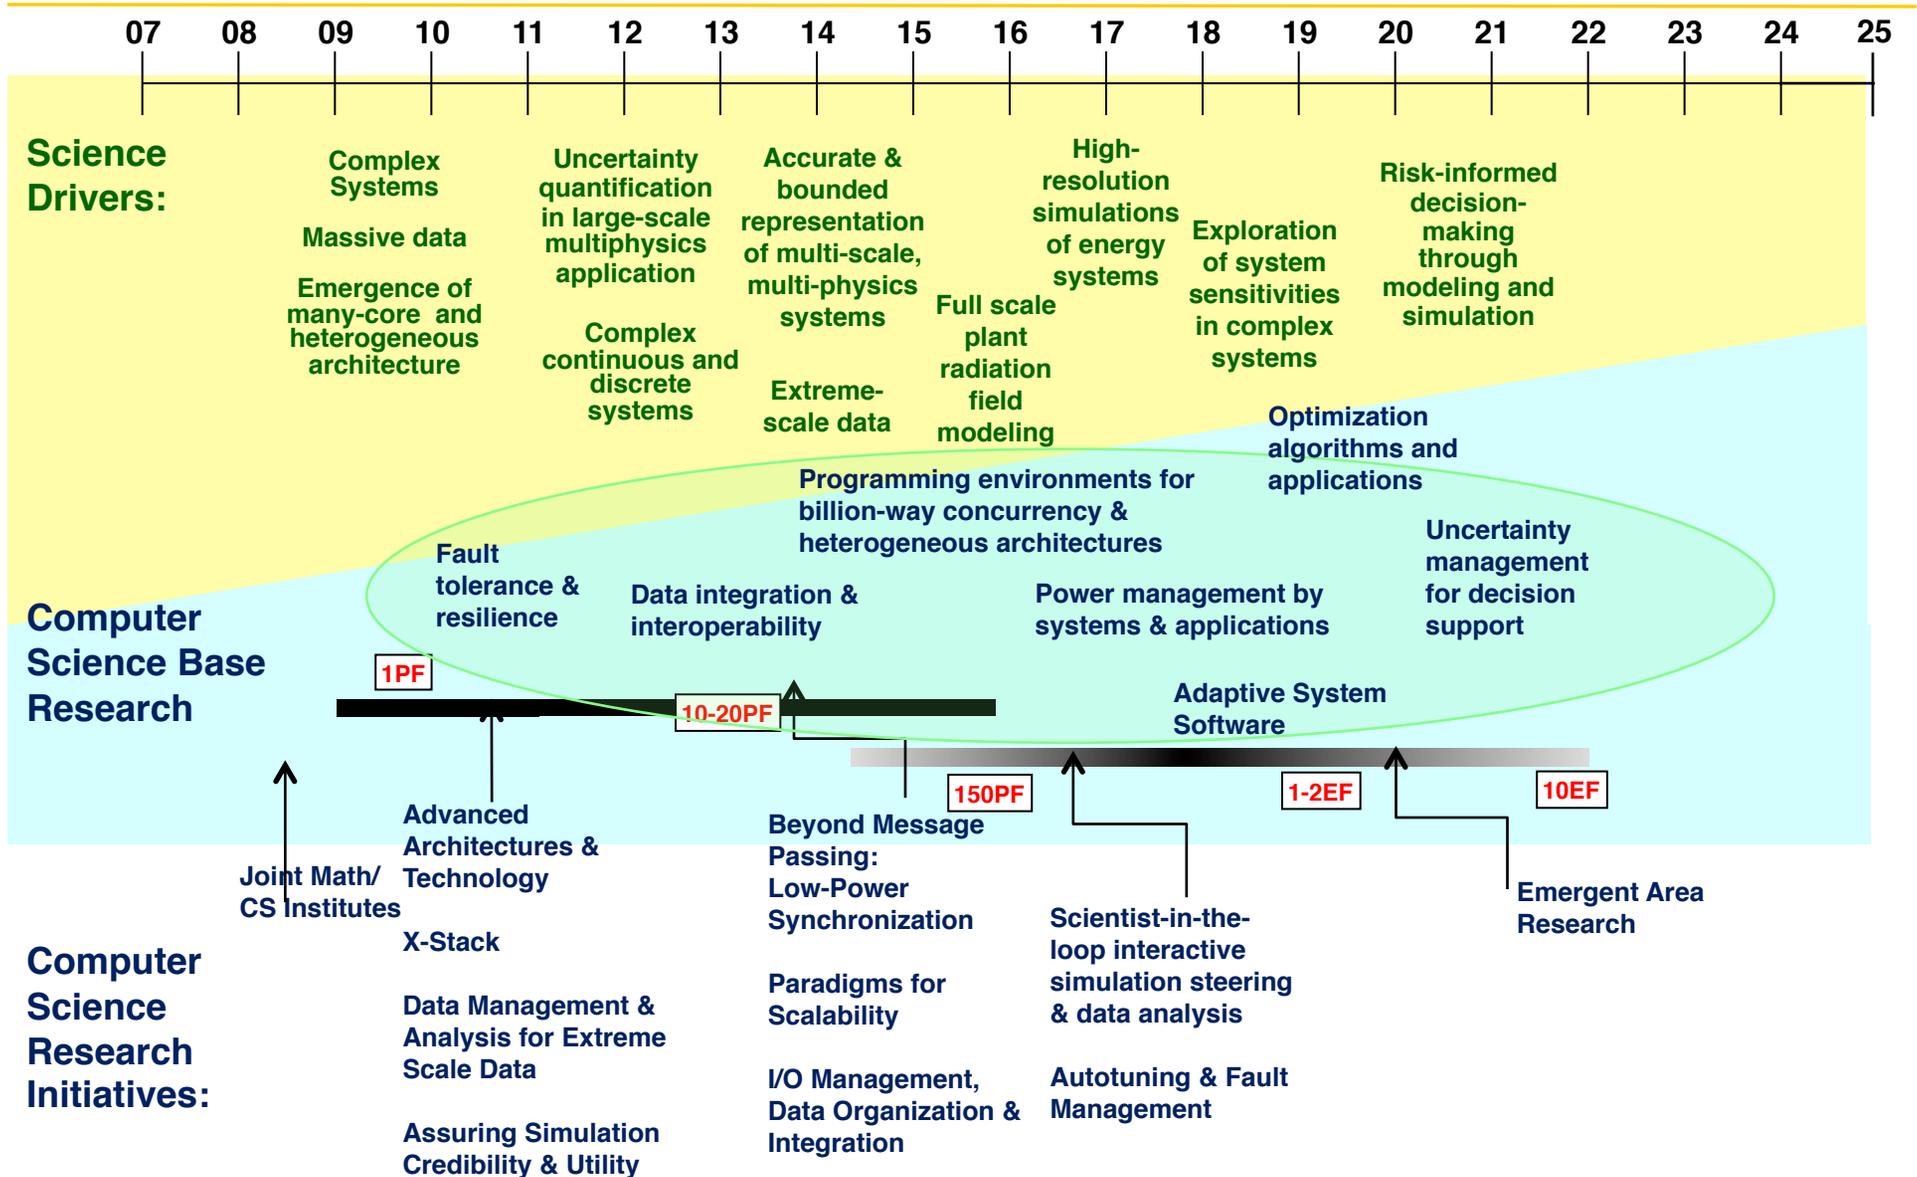


Application-Algorithm Challenges:
~10M-way concurrency





Computer Science Research: Creating Usable Exascale Environments for Computational Science





Next-Generation Network and Large-Scale Scientific Collaboration Research

